

# Ground Delay Program Analytics with Behavioral Cloning and Inverse Reinforcement Learning

Michael Bloem\*

*NASA Ames Research Center, Moffett Field, CA, 94035, USA*

Nicholas Bambos†

*Stanford University, Stanford, CA, 94305, USA*

We used historical data to build two types of model that predict Ground Delay Program implementation and also produce insights into how and why those implementation decisions are made. More specifically, we built behavioral cloning and inverse reinforcement learning models that predict hourly Ground Delay Program implementation at Newark Liberty International and San Francisco International airports. Data available to the models include actual and scheduled air traffic metrics and observed and forecasted weather conditions. We found that the random forest behavioral cloning models we developed are substantially better at predicting hourly Ground Delay Program implementation for these airports than the inverse reinforcement learning models we developed. However, all of the models struggle to predict the initialization and cancellation of Ground Delay Programs, which are both rare events. We also investigated the structure of the models in order to gain insights into Ground Delay Program implementation behavior. Notably, characteristics of both types of model suggest that GDP implementation decisions are made primarily based on conditions now or conditions anticipated in the next couple of hours.

## Nomenclature

$\mathcal{A}$	set of all GDP implemented actions (GDP and no GDP)
$a_t$	GDP implemented action in $t$ (GDP or no GDP)
$\text{AAR}_t$	airport arrival rate in $t$
$b_t^a$	number of flights in air buffer at $t$
$b_t^g$	number of flights in ground buffer at $t$
$\text{controlled}_t$	true if the arrival rate is controlled by a GDP in $t$ and false otherwise
$f(s, a)$	$M \times 1$ vector of reward features for state-action pair $(s, a)$
$\text{GDP}_t$	true if a GDP is implemented in $t$ and false otherwise
$M$	number of reward features
$q(s, a)$	score for state-action pair $(s, a)$ ; used by $\pi_C$
$R(s, a)$	reward for using action $a$ in state $s$
$R_C(s, a)$	reward relative to $\pi_C$
$\hat{R}_C(s, a)$	regressor that approximates $R_C(s, a)$
$\text{rate}_t$	controlled arrival rate for a GDP in $t$
$S$	set of all states
$s_t$	system state in $t$
$\text{scheduled}_t$	scheduled number of arrivals in $t$
$t$	time step

\*Research Aerospace Engineer, Systems Modeling and Optimization Branch, MS 210-15; member, AIAA; michael.bloem@nasa.gov.

†Professor of Management Science & Engineering and Electrical Engineering; bambos@stanford.edu

$\gamma$	discount factor
$\pi_C$	deterministic score function-based multi-class classifier policy
$\hat{\pi}_{\hat{R}_C}^*$	rollouts policy that approximately optimizes with respect to reward function $\hat{R}_C$
$\hat{\theta}$	$M \times 1$ vector of estimates of parameters in $\hat{R}_C$

## I. Introduction

WHEN predictions of air traffic and capacity suggest that an excessive number of flights will arrive at an airport at some future time, air traffic flow management (TFM) actions like Ground Delay Programs (GDPs) can be used to delay flights on the ground, where delay is relatively inexpensive. These actions are often used when a weather event such as high winds or low ceilings reduces the capacity of an airport. GDPs are used hundreds of times per year at some airports,<sup>1</sup> and each GDP can generate thousands of minutes of ground delay. TFM actions are selected by human decision makers who must rely primarily on experience and intuition because available forecasts of traffic and weather do not adequately account for uncertainties and because they have access to few “what-if” simulation capabilities or other decision-support tools.<sup>2,3</sup> In this research, we seek to use historical data to build models that can predict GDP actions and provide insight into how and why they are selected. These predictions and insights might be useful in the development of decision-support tools, for the training of air traffic managers, or when simulating airspace systems.

To gain insight into the decisions of these expert human decision makers and also to predict TFM actions for use in decision-support tools or simulations, researchers have begun applying *imitation learning* techniques in which demonstrations of the expert behavior found in historical data are used to develop models that mimic expert actions.<sup>4</sup> Some studies have used traditional classification or clustering techniques to build models that predict or suggest TFM actions based on features describing the state of the air traffic system.<sup>5-9</sup> For example, Mukherjee, Grabbe, and Sridhar compared the performance of two classification models (logistic regression and decision tree) that they trained to predict the probability of GDP implementation at an airport based on features describing the current weather and traffic state at the airport.<sup>9</sup> This approach to imitation learning is known as *behavioral cloning* (BC), and it assumes that expert actions can be characterized and modeled as a reaction to the current state. An alternative imitation learning approach is *inverse reinforcement learning* (IRL). IRL is based upon a model of the system dynamics and how these dynamics are affected by actions; a Markov decision process (MDP) model is typical. IRL assumes that the expert is selecting actions in an attempt to maximize a total reward accumulated over time while operating within the system model. IRL uses the demonstrations of the expert behavior found in historical data to infer a reward function that is consistent with the expert behavior (assuming strategic and total-reward-maximizing expert behavior). While BC can leverage powerful classification and clustering algorithms to make use of many features describing the system state, it does not consider system dynamics as explicitly as IRL.<sup>4</sup> Furthermore, in problems described as “long-range” and “goal-directed,” IRL has been shown to produce models that generalize to new environments better than models produced by BC.<sup>10</sup> All TFM problems certainly involve a dynamical system and, since GDP actions can cause ground delays for flights not scheduled to arrive at a constrained airport for several hours, seem to be strategic. Although we are only aware of applications of IRL to systems considerably simpler than the selection of GDP actions, these characteristics of GDPs suggest that IRL may produce models that can predict and provide insight into expert GDP actions. Even if IRL models produce relatively poor predictions, the reward function inferred by IRL algorithms may provide a “succinct, robust, and transferable” definition of when to implement GDPs.<sup>11</sup> Insights gleaned from such a reward function could guide researchers as they develop TFM decision-support tools.<sup>12,13</sup>

In this paper, we deploy both BC and IRL techniques to model and predict hourly expert GDP implementation actions. GDPs at Newark Liberty International (EWR) and San Francisco International (SFO) airports are analyzed. The two classes of imitation learning techniques will be provided with training data sets including historical scheduled and actual air traffic and forecasted and observed weather conditions as well as corresponding historical GDP actions. The models produced by the two classes of techniques will be compared based on their predictive performance when deployed on testing data sets and also based on the insights they provide.

The remainder of this paper is structured as follows. Section II reviews the data we use in this analysis. Next, we specify the imitation learning GDP models that are developed in this research in Section III. We evaluate the quality of model predictions with some experiments described in Section IV. We also investigate the models to glean insight into GDP implementation behavior in that Section. Finally, we

provide conclusions in Section V and suggest future work in Section VI.

## II. Data

In this Section, we describe the data that is used to quantify the system state for GDP models attempting to predict GDP actions. These data have been collected for EWR and SFO for the 151 days from 1 May 2011 through 29 September 2011, the 152 days from 1 May 2012 through 30 September 2012, and the 152 days from 1 May 2013 through 30 September 2013; there were 455 total days in the data set from these three summers. We studied EWR and SFO because in recent years, GDPs have been used at these airports more than at any other airports.<sup>1</sup> Indeed, there were 208 and 297 GDPs utilized at EWR and SFO, respectively, during these 455 days. Assuming no more than one GDP was used per airport per day (as is typical), a GDP was utilized at EWR in more than 45% of these days and at SFO in more than 65% of these days. Furthermore, we expect there to be meaningful differences in GDP decision-making practice for the two airports. This expectation is motivated in part by differences between the weather phenomena typically leading to decreased arrival capacities at these two airports (winds at EWR but low ceilings at SFO).<sup>1</sup> This expectation is also motivated by differences between the characteristics of arrival traffic demand at the two airports. For example, due to its relative proximity to multiple metropolitan areas in the eastern and midwestern United States, a larger fraction of flights bound to EWR than to SFO originate within a relatively short distance of the airport.

The causes of reduced airport capacity can vary between summer and winter months, so we simplified the problem by choosing to study only summer months. For example, some GDPs at EWR are reported as caused by “snow/ice,” but by using only summer months we avoid the need to develop models that can account for such GDPs.<sup>1</sup> Night-time data were removed from consideration because, due to low traffic volumes, GDPs are typically not used during the night. Hourly samples of each type of data were generated from 11:00 UTC (7:00 am EDT) through 06:00 UTC (2:00 am EDT) on the next day (20 hours per day) for EWR and from 12:00 UTC (5:00 am PDT) through 09:00 (2:00 am PDT) on the next day (22 hours per day) for SFO. Overall, this leads to 9100 hourly data points for EWR and 10,010 for SFO.

At each hour, some predictions of future states are available. For example, arrival schedules are a simple prediction of future arrival traffic levels, and weather forecasts provide predictions of future weather conditions. The set of features for each hour includes these predictions at one-hour intervals starting an hour from the current time and extending to the hour starting four hours from the current time. We chose to provide predictions that extend this far into the future because providing models with predictions extending further into the future did not improve predictive performance. Overall, 122 features describing the state were made available to GDP models. These features will be described in the following sub-sections.

### A. Weather Observations

Observations of weather conditions at an airport are recorded in METAR reports, which we retrieve from the FAA’s Aviation System Performance Metrics (ASPM) database.<sup>14</sup> These data include the meteorological conditions (instrument or visual), ceiling height and its change since the last hour, visibility distance, wind speed, and wind angle (six features).

### B. Weather Forecast

Terminal Aerodrome Forecasts (TAFs) provide predictions of weather conditions at an airport that extend 24 hours or more into the future, and they are often used by TFM decision makers. At EWR and SFO, TAFs were published at least every three hours during the time period we are studying. From TAFs, we extract predictions of conditions at the hour starting at the current time through the hour starting four hours from the current time. For each of these five hours, we include a feature for the meteorological conditions (visual, marginal visual, instrument, or low instrument), ceiling height and its change since the last hour (change only computed for hours after the current hour), whether ceiling heights are “temporary,” wind speed, wind direction, wind gust speed, landing runway cross wind speed for the preferred runway configuration, landing runway head wind speed for the preferred runway configuration, visibility distance, whether the visibility distance is greater than the specified quantity, whether the visibility conditions are “temporary,” the intensity of precipitation, the intensity of obscuration, and whether precipitation and obscuration are “temporary.”

There are 14 features specified for 5 hours and one (change in ceiling height) for 4 hours, which means there are 74 features derived from the TAF for each hour.

### C. Number of Scheduled Arrivals

ASPM records also contain the number of scheduled arrivals at the airport during each quarter hour. We use these records to generate a feature denoting the scheduled arrivals during the hour starting at the current time and extending through the hour starting four hours from now (five features).

### D. Current Airport State

Two other features that describe the current state of the airport are also extracted from ASPM records: the airport arrival rate (AAR) for the current hour and the runway configuration.

### E. Predictions of Future Airport Arrival Rates

GDPs are typically used when an airport's arrival capacity, quantified by the AAR, is expected to be too low to handle the predicted number of arrivals. The AAR selected by traffic managers depends on factors like runway configurations, weather conditions, and the type of aircraft that are scheduled to arrive at the airport. Other researchers have developed AAR prediction models,<sup>15–22</sup> and we implemented a model similar to the bagged decision tree model proposed by Provan, Cunningham, and Cook in Refs. 20 and 21. This type of model worked well not only for Provan, Cunningham, and Cook, but also for Wang in Refs. 17 and 19. The AAR predictions from this model for the hour starting one hour from the current time through the hour starting four hours from the current time are provided to the GDP models (four features).

### F. Reroutes

Reroute advisories were collected from the FAA's National Traffic Management Log (NTML) and used to construct features describing reroutes required or recommended for flights bound to the airport. GDPs are sometimes used to help address reductions in the capacity of airspace typically used by flights headed to the airport, such as airspace near arrival fixes. This sort of capacity reduction also may result in reroutes, so these features attempt to quantify reductions in airspace capacity that may cause GDP implementation. More precisely, for the hour starting at the current time through the hour starting 4 hours from the current time, there are 4 reroute-related features (20 total). Two of these report the number of departure Air Route Traffic Control Centers (ARTCCs or Centers) and departure airports for which reroutes are recommended for flights bound to the airport. The other two features report the number of departure Centers and departure airports for which reroutes are required for flights bound to the airport.

### G. Previous GDP Plan

GDP plan data was also retrieved from the FAA's NTML. While GDP plans are sometimes modified, selected GDP actions often are a continuation of a previously-announced plan. For each hour, we quantify the GDP action that would be pursued at the start of the hour, assuming the GDP plan as specified one hour ago is simply continued. The features describing the previous plan include whether or not there would be a GDP implemented, the GDP scope<sup>a</sup>, the number of hours until the first hour of controlled GDP rates, the number of hours until the last hour of controlled GDP rates, and the GDP rate that should be in place for each relevant hour of the GDP (from the hour starting at the current time through the hour starting four hours from the current time). This means that there are nine features describing the previous GDP plan for each hourly time step.

### H. Ground and Air Buffers

One of the fundamental purposes of a GDP is to prescribe ground delay, which is cheaper than delay absorbed in the air, when some delay must be absorbed. Therefore, we defined a simple deterministic queuing network model with buffers for flights delayed on the ground and in the air. It provides estimates of two quantities

---

<sup>a</sup>The scope is a distance typically measured in miles or a number of Air Route Traffic Control Centers from the destination airport. Flights departing from further than the specified scope are not subject to ground delays imposed by a GDP.

that are essential to GDP planning: how much delay is absorbed on the ground and in the air as a result of GDP actions. This model is similar to the deterministic queuing model utilized by Kim and Hansen in Ref. 23 and the stochastic queuing model proposed by Odoni and discussed by Ball et al. in sub-section 4.3 of Ref. 24. Although this model certainly fails to capture several relevant aspects of GDPs (such as differences in flight times between flights), our hope is that it is simple but not simplistic: that it captures enough of the relevant characteristics of the real world to be useful for GDP analytics, but without the burden of unnecessary complexity.

The number of flights in the ground buffer at the start of time step  $t$  is  $b_t^g$  and the number in the air buffer is  $b_t^a$ . At the start of each day, the ground and air buffers are initialized to zero. Then, the system dynamics specify that the buffer levels at the start of the next time step ( $b_{t+1}^g$  and  $b_{t+1}^a$ ) depend on the scheduled arrivals ( $\text{scheduled}_t$ ) and AAR ( $\text{AAR}_t$ ) during time step  $t$ , as well as  $b_t^g$ ,  $b_t^a$ , and the GDP action implemented in this time step. If a GDP was implemented during  $t$  ( $\text{GDP}_t$ ), then the relevant components of the GDP action are whether or not the arrival rate is controlled by the GDP during  $t$  ( $\text{controlled}_t$ ) and the controlled arrival rate during  $t$  ( $\text{rate}_t$ ). The buffer levels are updated according to

$$b_{t+1}^g = \begin{cases} [b_t^g + \text{scheduled}_t - \text{rate}_t]_+ & \text{if } \text{GDP}_t \text{ and } \text{controlled}_t \\ 0 & \text{else} \end{cases} \quad (1)$$

and

$$b_{t+1}^a = \begin{cases} [b_t^a + \min(b_t^g + \text{scheduled}_t, \text{rate}_t) - \text{AAR}_t]_+ & \text{if } \text{GDP}_t \text{ and } \text{controlled}_t \\ [b_t^a + b_t^g + \text{scheduled}_t - \text{AAR}_t]_+ & \text{else.} \end{cases} \quad (2)$$

Here  $[x]_+$  is equal to  $x$  if  $x \geq 0$  but equal to 0 if  $x < 0$ . Fig. 1 depicts this simple queuing model when a GDP is implemented.



Figure 1: Ground and air buffer system model.

### III. GDP Models

The structure we specified and used for both BC and IRL GDP models is depicted in Fig. 2. The input data that can be used by the models is described in Section II. The output of the model is either a prediction that a GDP would not be implemented in the state quantified by the input data, or a prediction that a GDP would be implemented, along with predictions of the GDP parameters that would be used in the state quantified by the input data. The GDP parameters include the GDP scope, the time when the controlled rates in the GDP begin (the GDP *start time*), the number of hours of enforced GDP rates (the GDP *duration*), and the enforced rate for each hour in the GDP duration, extending from the hour starting at the current time through the hour starting four hours in the future (up to five hourly rates). There are two sub-models: one that predicts whether or not a GDP will be implemented and another that predicts GDP parameters when a GDP is initialized. The model is a simplification of reality in that it requires that a GDP plan either progress as planned or be canceled (no modifications or extensions are permitted).

#### A. GDP Implemented Models

A GDP implemented model predicts whether or not a GDP will be implemented during a given hour, given a set of features describing the state (see Section II). BC and IRL GDP implemented models will be developed and compared. The specific BC and IRL algorithms will be discussed in the next two sub-sections.

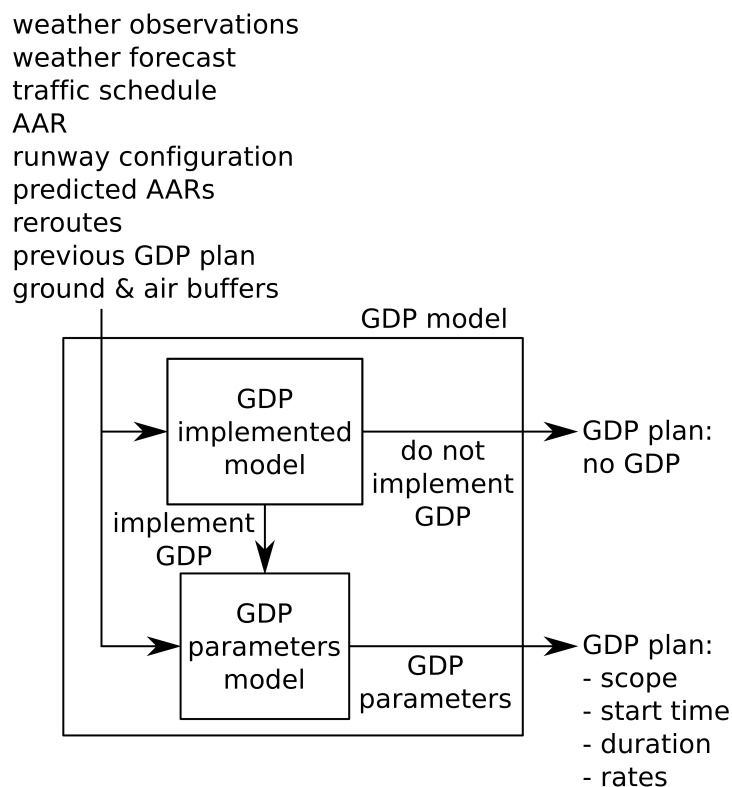


Figure 2: Structure of GDP model.

### 1. BC: Random Forest Classifiers for Cancellation and Initialization

The structure of the BC model developed and analyzed for this research is depicted in Fig. 3. Depending on whether or not the previous GDP plan specified that a GDP would be implemented in the current hour, either a GDP cancellation or GDP initialization model is invoked. We hope that creating specific models for what seem to be different decisions (GDP cancellation and GDP initialization) will lead to better predictive performance, as well as to more refined insight into GDP decision making. An additional motivation for building separate models for GDP cancellation and GDP initialization is that doing so facilitates the use of over- and under-sampling to generate custom training data sets for each model, which helps each with the difficult imbalanced classification problem it faces.

The GDP initialization model is provided with features describing the current state at the airport, but no features describing the previous GDP plan because the previous plan was to not use a GDP. It then predicts either that a GDP will be initialized or that no GDP will be initialized. The GDP cancellation model is provided with features describing not only the current state at the airport, but also features describing the previous GDP plan, such as the scope, planned end time, and rates. It then predicts either that the GDP will continue as planned or that it will be canceled. In hours that are immediately after the planned end of a GDP, the GDP cancellation model is *not* used because not implementing a GDP in that hour is *not* a cancellation. The question in those hours is whether or not a new GDP will be initialized, so the GDP initialization model is invoked.

The GDP cancellation and GDP initialization models are both random forest classification models, implemented with the `RandomForestClassifier` class available in the scikit-learn package for the Python programming language.<sup>25</sup> Random forest models were selected because they typically perform well with minimal tuning of the algorithms that train the models, they generally do not over-fit to the training data even when provided with many features, and other researchers have found that related models predict AARs better than alternative models.<sup>17, 19, 20, 26</sup>

Among all of the hours where a GDP is planned to be implemented, only in about one in ten hours is a



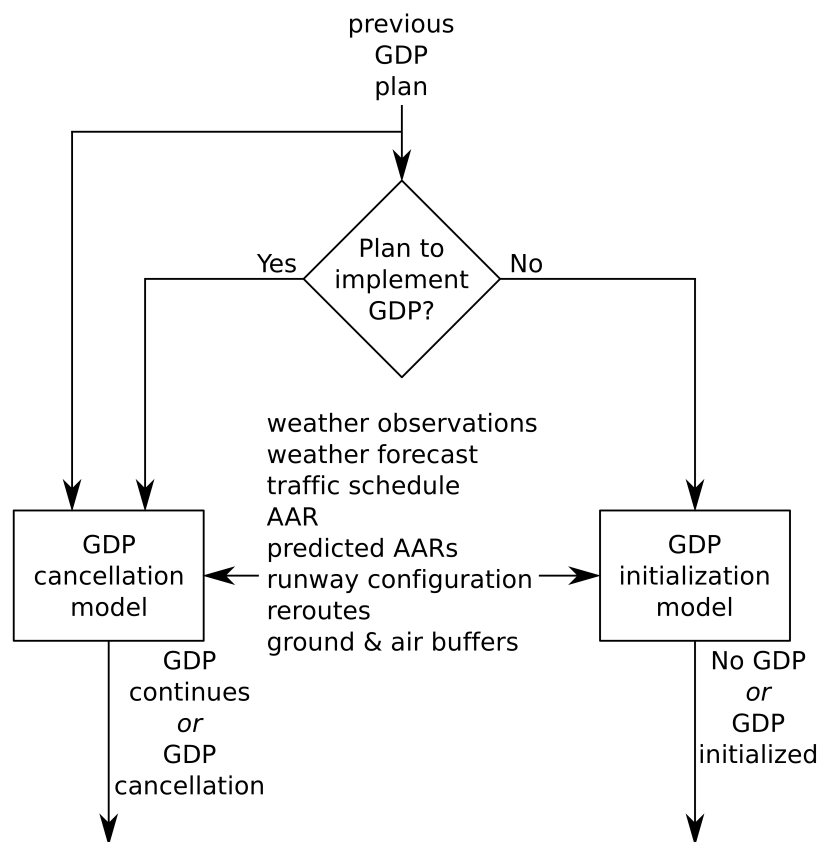


Figure 3: Structure of the BC GDP implemented model.

GDP canceled in the data we are analyzing. Similarly, among all the hours where a GDP is not planned to be implemented, only in about one in thirty hours is one initialized. When facing imbalanced data sets such as these, predictive performance can sometimes be enhanced by using the Synthetic Minority Over-Sampling Technique (SMOTE).<sup>27</sup> We used an implementation of the SMOTE algorithm to generate synthetic minority class (initialization and cancellation) data points,<sup>28</sup> and we also under-sampled the majority class data points. In particular, the number of minority samples was doubled by using SMOTE to generate synthetic samples, and the majority samples were under-sampled so that there were four times as many (for EWR) or twice as many (for SFO) majority-class samples as (real and synthetic) minority-class samples.

## 2. IRL: Cascaded Supervised Inverse Reinforcement Learning

The IRL algorithm selected for evaluation is the Cascaded Supervised Inverse Reinforcement Learning (CSI) algorithm proposed by Klein, Piot, Geist, and Pietquin in Ref. 29. This approach was selected for several reasons: it is relatively easy to implement, it does not require multiple computations of an optimal policy for various possible reward functions, and it leverages existing classification and regression algorithms. Furthermore, it does not involve exploring the entire state space—just the states visited in the training data and possibly in some additional simulations. Finally, the expert policy that produced the demonstrations in the data is near-optimal for the reward function estimated by the CSI algorithm (see Ref. 29, Theorem 1). One downside of the algorithm is that it assumes a deterministic expert policy that is optimal for a certain reward function, which may not be the case for GDP decision making. This drawback did not prevent us from selecting the CSI algorithm, however, because for this initial investigation we are working with deterministic GDP implemented models.

There are six main elements involved in CSI. Each of these elements will be described in the subsequent paragraphs.

*System Model:* The CSI algorithm, like any other IRL algorithm, requires a system model. The system state at a time step  $t$  is  $s_t$ , which is a member of a set of all possible states  $\mathcal{S}$ . It involves an *exogenous* state and a *controlled* state. The exogenous state is a vector of features describing the weather observations, the weather forecast, the traffic schedule, the current AAR, the runway configuration, predicted AARs, and reroutes (see Section II). GDP actions have no impact on the exogenous state dynamics in this model, and the CSI algorithm allows us to proceed even though we assume that the state transition probabilities for this part of the state are not specified. The controlled state contains the GDP action for this time step prescribed by the previous GDP plan,  $b_t^a$ , and  $b_t^g$ . The controlled state dynamics are impacted by GDP actions as prescribed by Eqs. (1) and (2). The action  $a_t$  taken at  $t$  is binary: it specifies either that no GDP is implemented or that a GDP is implemented in this time step ( $a_t \in \mathcal{A} = \{\text{GDP, no GDP}\}$ ). This is a Markov model because the conditional distribution of future states depends only on the current state and action, not on the whole history of states and actions.

*Decision Process Model:* We assume that traffic managers are attempting to maximize the expected value of a discounted infinite sum of future rewards. More precisely, they seek a deterministic policy mapping states to actions that maximizes  $\mathbf{E}[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t)]$ , where  $\gamma \in (0, 1)$  is the discount factor and  $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbf{R}$  is the reward for each time step. We choose an appropriate value for  $\gamma$  by selecting from a set of possible values the one that achieves the best fit (as measured by  $R^2$  on testing data) when we conduct the reward estimation step later in this sub-section. This value was 0.4 for EWR. For SFO, this value was actually 0.0 but we utilized a value of 0.2 instead because the fit was not much worse at  $\gamma = 0.2$  and because  $\gamma = 0.0$  suggests no strategic expert behavior, making the IRL algorithm simply a complex BC model.

*Classifier Policy and Estimation of Optimal State-Action Value:* Once the system and decision process model have been specified, the next step in CSI is to derive a deterministic classifier policy  $\pi_C : \mathcal{S} \rightarrow \mathcal{A}$ . This classifier policy can be any *score function-based multi-class classifier* (SFMC<sup>2</sup>). An SFMC<sup>2</sup> predicts an action that achieves the largest score according to some score function  $q(s, a)$ :  $\pi_C(s) \in \arg\max_{a \in \mathcal{A}} q(s, a)$ . By interpreting the score function  $q(s, a)$  as an optimal *state-action value* function for  $\pi_C$  and making use of the Bellman equation, the CSI algorithm can use the score function and the policy to generate a reward sample data point corresponding to each state-action pair. We will utilize a model similar to the random forest BC GDP implemented model described in sub-section III.A.1 for the classifier policy  $\pi_C$ . The score function for a given state and action is the average over all the trees in the random forest of the fraction of members in the leaf nodes for which the action was selected.

*Reward Estimation with Regression:* The next step in CSI is to use a regression algorithm to estimate a reward function  $R_C$  that is consistent with the reward samples produced by using the Bellman equation and the score function  $q$  when it is interpreted as an optimal state-action value function for  $\pi_C$ .<sup>29</sup> Any type of regressor can be used, but we use a linear regressor model. Linear models are relatively easy to interpret, so a trained linear regressor model should be relatively rich in insight. Furthermore, some proposed TFM and GDP optimization approaches assume or even require a reward function that is linear in the optimization decision variables (Refs. 2 and 24 describe examples of such approaches).

The form of a linear regressor model is  $\hat{R}_C(s, a) = \hat{\theta}^\top f(s, a)$ , where  $f(s, a) \in \mathbf{R}^M$  is a vector of reward features for the state-action pair  $(s, a)$  and  $\hat{\theta} \in \mathbf{R}^M$  is a vector of parameters estimated by the regression algorithm. The reward features we utilize were inspired by the objective function used in the Ground Delay Program Parameter Selection Model<sup>12</sup> and also by performance measures suggested by Ball et al.<sup>24</sup> The eight reward features are ground and air buffer levels at the end of the current time step, the change in the air and ground buffers during this time step, an indicator that the air buffer will be greater than or equal to five at the end of the current time step, the number of arrivals during the time step, the number of unused arrival slots while the arrival rate is controlled by a GDP (i.e. while  $\text{controlled}_t$  is true), and the canceled duration of a canceled GDP.

*Derivation of Approximately-Optimal Policy:* We derive a policy  $\hat{\pi}_{R_C}^*$  that attempts to maximize the infinite discounted total reward objective, defined based on  $\hat{R}_C$ , using an approximate dynamic programming approach known as *rollouts* (see Ref. 30, sub-section 6.4). Ultimately, the prediction produced by the CSI model of whether or not a GDP will be implemented when in state  $s$  will be the action returned by  $\hat{\pi}_{R_C}^*(s)$ . Other techniques from reinforcement learning and approximate dynamic programming could be used to derive this policy. We selected rollouts because we could use weather forecast and traffic schedule data already in the system state to perform the simulations required when estimating optimal state-action value functions, which made the problem tractable in spite of the unknown state transition probabilities.



## B. GDP Parameters Model

As depicted in Fig. 2, a GDP parameters model was developed to predict the GDP scope, GDP start time, GDP duration, and GDP rates. This model is only used in the final policy estimation step of the CSI algorithm, described in sub-section III.A.2. Random forest BC models were used for each sub-model that predicts one of the GDP parameters, and the predictions were rounded to the nearest typical value for the parameter in question. Random forest models were selected because they perform well with relatively little tuning of the algorithm that trains the models, they generally do not over-fit to the training data even when provided with many features, and other researchers have had some success using related models to predict AARs.<sup>17, 19, 20, 26</sup> Random forest regression models for these parameters were trained with the `RandomForestRegressor` class in the scikit-learn Python package<sup>25</sup> using settings and parameters similar to those suggested in Ref. 20. Although these models are not the focus of this research, we are not aware of previous attempts at building models that predict GDP parameters.

## IV. Experiments

Ten-fold cross validation is used in these experiments (see Ref. 26, sub-section 7.10.1). The folds are defined based on days in the data set rather than individual time steps (hours). With 455 days in the data set and ten folds, each fold consisted of all of the time steps in around 46 days. The experiments focus on the GDP implemented models.

### A. Prediction Quality Results

The prediction quality metrics for GDP implemented models are computed based on three confusion matrices that can be constructed with the predictions for the testing data. The three matrices investigated here are constructed with the data in the ten test data folds. Each hour-long time step in the full data set is in a testing data fold exactly once, so the testing data contains each sample from the full data set exactly once. The first confusion matrix describes predictions of hourly GDP implementation, the second describes predictions of GDP initialization, and the third describes predictions of GDP cancellation. The first matrix involves all the testing data points, the second and third matrices are based on only some of the testing data points, and each data point is involved in either the second or the third matrix but not both. For each confusion matrix, the accuracy, precision, recall, and  $F_1$ -score will be reported. Each of these metrics will be in the range  $[0, 1]$ , with larger values indicating better predictive performance. Since precision and recall are particularly relevant for imbalanced data sets, such as those faced by the initialization and cancellation models, and since the  $F_1$ -score is the harmonic mean of these two metrics, we view the  $F_1$ -score as the most important single metric that quantifies the predictive performance for each confusion matrix.

#### 1. Baseline: Quality of GDP Plan Model Predictions

We will summarize the quality of predictions produced by a baseline model that simply predicts that the previous GDP plan will be executed. If no GDP is planned for the hour in question, the model predicts that no GDP will be initialized, and if a GDP is planned, then the model predicts that it will continue and not be canceled. The  $F_1$ -scores achieved by this model are 0.90 and 0.87 for GDP implementation at EWR and SFO, respectively, which is suggestive of high-quality predictions. However, this model achieves a recall of 0.00 and an undefined precision and  $F_1$ -score for predictions of initializations and cancellations at both airports. Initialization and cancellation events are important operationally, so this model's failure to predict these events reveals its limited operational value. Furthermore, this model's performance illustrates the importance of evaluating models based on their predictions of initializations and cancellations, not just their predictions of GDP implementation.

#### 2. Quality of BC GDP Implemented Model Predictions

Tables 1 and 2 show the three confusion matrices and related metrics achieved by the BC GDP implemented models for EWR and SFO, respectively, when they are presented with the ten test data folds. For both EWR and SFO, the accuracy, precision, recall, and  $F_1$ -score of the overall GDP implemented models are relatively close to one. However, this strong overall performance masks how much the GDP initialization and GDP cancellation models struggle to predict the rare initialization and cancellation events. Although the  $F_1$ -scores

for predictions of GDP initialization and cancellation are slightly higher for SFO than for EWR, they all range between 0.41 and 0.63, which is low. For both EWR and to a lesser extent SFO, the low precision scores achieved by the GDP initialization models indicate a high false alarm rate—that they predict that GDPs will be initialized much more frequently than they actually are initialized. The GDP initialization data set is particularly imbalanced, with 34 and 26 times more non-initialization events than initialization events for EWR and SFO, respectively. This makes predicting initializations difficult. The GDP cancellation models also suffer from low precision scores, and again this is partially a result of the imbalanced nature of the data set. There are 11 and 7 times more non-cancellation events than cancellation events for EWR and SFO, respectively.

Table 1: Confusion Matrices for EWR BC GDP Implemented Model

(a) *Implementation*: accuracy=0.94, precision=0.83, recall=0.92,  $F_1$ -score=0.88

Actual	Predicted No GDP	Predicted GDP	Total
No GDP	6801	357	7158
GDP	151	1791	1942
Total	6952	2148	9100

(b) *Initialization*: accuracy=0.95, precision=0.31, recall=0.62,  $F_1$ -score=0.41

Actual	Predicted No Initialization	Predicted Initialization	Total
No Initialization	6707	286	6993
Initialization	80	128	208
Total	6787	414	7201

(c) *Cancellation*: accuracy=0.93, precision=0.57, recall=0.57,  $F_1$ -score=0.57

Actual	Predicted Continuation	Predicted Cancellation	Total
Continuation	1663	71	1734
Cancellation	71	94	165
Total	1734	165	1899

Table 2: Confusion Matrices for SFO BC GDP Implemented Model

(a) *Implementation*: accuracy=0.95, precision=0.84, recall=0.90,  $F_1$ -score=0.87

Actual	Predicted No GDP	Predicted GDP	Total
No GDP	7616	346	7962
GDP	203	1845	2048
Total	7819	2191	10010

(b) *Initialization*: accuracy=0.96, precision=0.49, recall=0.88,  $F_1$ -score=0.63

Actual	Predicted No Initialization	Predicted Initialization	Total
No Initialization	7443	270	7713
Initialization	36	261	297
Total	7479	531	8010

(c) *Cancellation*: accuracy=0.88, precision=0.51, recall=0.69,  $F_1$ -score=0.59

Actual	Predicted Continuation	Predicted Cancellation	Total
Continuation	1584	167	1751
Cancellation	76	173	249
Total	1660	340	2000

### 3. Quality of IRL GDP Implemented Model Predictions

Tables 3 and 4 show the three confusion matrices and related metrics achieved by the IRL GDP implemented models for EWR and SFO, respectively, when they are presented with the ten test data folds. For both EWR and SFO, the IRL GDP implemented models demonstrate substantially lower predictive performance than the BC GDP implemented models. This is most evident for predictions of GDP initialization: the EWR IRL model predicts initialization much too frequently while the SFO IRL model predicts initialization too infrequently. The SFO IRL model also predicts GDP cancellation too infrequently.

Table 3: Confusion Matrices for EWR IRL GDP Implemented Model

(a) *Implementation*: accuracy=0.80, precision=0.52, recall=0.92,  $F_1$ -score=0.67

Actual	Predicted No GDP	Predicted GDP	Total
No GDP	5535	1623	7158
GDP	159	1783	1942
Total	5694	3406	9100

(b) *Initialization*: accuracy=0.78, precision=0.084, recall=0.67,  $F_1$ -score=0.15

Actual	Predicted No Initialization	Predicted Initialization	Total
No Initialization	5468	1525	6993
Initialization	68	140	208
Total	5536	1665	7201

(c) *Cancellation*: accuracy=0.90, precision=0.42, recall=0.41,  $F_1$ -score=0.41

Actual	Predicted Continuation	Predicted Cancellation	Total
Continuation	1643	91	1734
Cancellation	98	67	165
Total	1741	158	1899

Table 4: Confusion Matrices for SFO IRL GDP Implemented Model

(a) *Implementation*: accuracy=0.94, precision=0.85, recall=0.84,  $F_1$ -score=0.84

Actual	Predicted No GDP	Predicted GDP	Total
No GDP	7647	315	7962
GDP	320	1728	2048
Total	7967	2043	10010

(b) *Initialization*: accuracy=0.95, precision=0.17, recall=0.088,  $F_1$ -score=0.12

Actual	Predicted No Initialization	Predicted Initialization	Total
No Initialization	7585	128	7713
Initialization	271	26	297
Total	7856	154	8010

(c) *Cancellation*: accuracy=0.88, precision=0.56, recall=0.25,  $F_1$ -score=0.34

Actual	Predicted Continuation	Predicted Cancellation	Total
Continuation	1702	49	1751
Cancellation	187	62	249
Total	1889	111	2000

Although it is difficult to determine exactly what causes the relatively poor predictive performance of the IRL GDP implemented models, the results of diagnostic tests<sup>31</sup> and analysis of the estimated regressor parameters (see sub-section IV.B.2) suggest that the poor predictive performance of the reward regressor is

likely an important cause. The reward regressor’s poor performance may in turn be the result of using invalid assumptions to specify the regression problem (such as CSI’s assumption of a deterministic and optimal expert policy), of selecting a reward regressor model form that is not suited to the regression problem, or of using an inaccurate or incomplete set of reward features. Several reward features are derived from the state of the ground and air buffer system model, and our intuition is that the fidelity of this system model may need to be improved.

## B. Insight Results

### 1. Insight from BC GDP Implemented Model

The main form of insight available from the random forest models used in the BC GDP implemented model is feature *importance scores*. The importance score for a feature is the total decrease in node “impurity” (as measured by the Gini splitting criterion) resulting from splits defined based on the feature and weighted by the proportion of samples reaching the corresponding nodes, averaged over all the trees in the ensemble. Based on this definition, important features define frequently-used splits and/or generate large improvements in the split criterion when they are used to define splits. Larger scores imply greater importance. We recorded the importance scores of the input features used by the models for each of the ten times the models are trained.

Figures 4 and 5 show the scores for the features with the ten highest importance scores for initialization and cancellation models that make up the BC GDP implemented models for EWR and SFO, respectively. The height of each bar is the mean of the importance scores for each feature according to the 10 models constructed for the ten training data sets used in cross validation, and the error bars show the standard deviation of these ten importance scores. Features names ending with a “ $k$ ” for some integer  $k$  are predictions of the feature for the hour starting  $k$  hours from the time of the prediction.

For the EWR GDP initialization model, features related to the AAR (“AAR”) or a prediction of the AAR (“Pred\_AAR”) make up five of these ten features. This suggests that future AAR levels have a relatively strong influence on GDP initialization, which makes sense given that GDPs are used when expected arrival demand exceeds expected arrival capacity and that arrival capacity is quantified by AAR. Features related to scheduled arrivals (“SCHARR”) account for another four of the ten features, which is also not surprising because GDPs are used to reduce arrivals when predictions suggest that there might be an excessive number of arrivals. The tenth feature in this set is a prediction of the number of departure Centers for which reroutes are required three hours in the future (“Centers\_Rerouted\_3”), suggesting that GDPs are initialized at EWR to help reduce demand for constrained airspace. This makes sense because flights bound for EWR typically traverse highly congested airspace in the northeastern United States.


For the EWR GDP cancellation model, five features related to parameters of the previous GDP plan, such as the planned time until the end of the GDP (“Prev\_GDP\_LATS\_to\_end”<sup>b</sup>) and planned GDP rates (“Prev\_GDP\_Rate”), achieve high average importance scores. This suggests that GDPs tend to be canceled when the previous plan indicates that they are close to finishing, which makes sense because presumably experts select appropriate end times when constructing the GDP plan. Features related to scheduled arrivals (“SCHARR”) account for four of the remaining features achieving the ten highest importance scores. This makes sense because if scheduled arrivals are low, a GDP may no longer be needed. The final feature in this set is a reroute-related feature (“Centers\_Rerouted\_2”), again suggesting that EWR GDPs might be partially caused by congested airspace that also requires new routes for flights bound for EWR.

Of the ten features with the largest average importance scores for the SFO GDP initialization model, four are related to observations or forecasts valid at some hour in the upcoming three hours of the ceiling or meteorological conditions (“CEILING”<sup>c</sup>, “Ceiling\_0”<sup>d</sup>, “Ceiling\_1”, and “Met\_Conds\_1”). This is not surprising given that SFO GDPs are largely caused by low ceilings.<sup>1</sup> Only one feature in this set is an AAR (“Pred\_AAR\_2”), which is fewer than the five such features for the EWR GDP initialization model. This might be because while the weather conditions that lead to lower capacity at SFO are relatively straightforward to describe and quantify, enabling the SFO model to identify and depend upon them directly, the capacity at EWR is more a more complicated function of a variety of weather and other conditions, leading


<sup>b</sup>“Prev\_GDP\_LATS\_to\_end” is the number of look-ahead time steps (hours) until the end of the GDP prescribed by the previous GDP plan.

<sup>c</sup>“CEILING” is the observation of the ceiling at the current hour recorded in a METAR report.

<sup>d</sup>“Ceiling\_0” is the prediction of the ceiling from a TAF forecast that is valid for the current hour. The TAF forecast may have been published very recently, or it may have been published up to three hours ago.



(a) GDP initialization model.




(b) GDP cancellation model.

Figure 4: Features with highest importance scores for the EWR BC GDP implemented model.


the EWR model to depend on AAR predictions (rather than directly on weather conditions). Features related to scheduled arrivals make up five of the ten features, which is also not surprising for reasons described earlier.

Finally, for the SFO GDP cancellation model, five of the ten features achieving the largest average importance scores come from the previous GDP plan (such as “Prev\_GDP\_LATS\_to\_end” and planned GDP rates). This result is consistent with the intuitive claim that the current GDP plan, particularly the planned remaining duration, is helpful when predicting SFO GDP cancellations. The planned rates may be important because they are another way to learn the planned remaining duration, or perhaps because they may indicate the degree to which capacity is diminished (GDPs might be less likely to be canceled when capacity is lower). Predicted or current AARs make up four other of these features, and the current ceiling is the final feature in the set. Unlike in the EWR GDP cancellation model, no features related to scheduled arrivals achieve top-ten importance scores in the SFO GDP cancellation model.

As was mentioned in Section II, we found that providing features describing predictions extending beyond four hours did not improve the predictive power of this type of model. This provides some insight: it is consistent with the claim that decision makers are not concerned with conditions extending beyond four



(a) GDP initialization model.



(b) GDP cancellation model.

Figure 5: Features with highest importance scores for the SFO BC GDP implemented model.

hours in the future when they determine whether or not to use a GDP.

## 2. Insight from IRL GDP Implemented Model

Although our evaluation suggests that the predictive power of the IRL GDP implemented model is low, the estimated parameters  $\hat{\theta}$  for the reward function regressor  $\hat{R}_C(s, a) = \hat{\theta}^\top f(s, a)$  still may provide some useful insights. Even these insights should be viewed with suspicion, however, because  $\hat{R}_C$  struggled to fit the reward sample training data. In the ten training data sets used in cross validation, the average  $R^2$  value achieved by  $\hat{R}_C$  was only 0.29 for EWR and 0.23 for SFO. These low values suggest that  $\hat{R}_C$  is not explaining much of the variation of the reward samples in the training data set. We suspect that this poor performance is an important cause of the poor predictive power of the IRL GDP implemented model, and improving the performance of the reward regressor is a topic for future research.

Table 5 shows average reward parameter estimates and corresponding average  $p$ -values for the ten regressors trained with the ten training data folds used in cross validation. Before training  $\hat{R}_C$ , scalar-valued reward features were standardized over the time steps for which they are defined, while the lone indicator feature (which takes a value of one when the air buffer at the end of the time step is greater than or equal



to five and zero otherwise) was not standardized. This facilitates interpretation of the parameter estimates by making comparisons of their relative magnitudes more meaningful. When constructing reward samples to train  $\hat{R}_C$ , we used  $\gamma = 0.4$  for EWR and  $\gamma = 0.2$  for SFO. Therefore, the ranges of possible and typical reward sample values differ between the two airports, and so we cannot directly compare the values of the estimated parameters between airports.

Table 5: Average Estimated Parameters and Average  $p$ -Values of  $\hat{R}_C$  for EWR and SFO

Reward Feature ( $f_m(s, a)$ )	EWR		SFO	
	$\hat{\theta}_m$	$p$ -value	$\hat{\theta}_m$	$p$ -value
constant	0.53	< 0.001	0.71	< 0.001
ground buffer at end of time step	0.0049	0.097	-0.017	< 0.001
change in ground buffer	0.0074	0.17	-0.00062	0.46
air buffer at end of time step	0.014	< 0.001	-0.0023	0.50
change in air buffer	-0.035	< 0.001	-0.035	< 0.001
air buffer at end of time step $\geq 5$	-0.24	< 0.001	-0.20	< 0.001
arrivals	-0.029	< 0.001	0.0071	0.011
unused slots during rate control	-0.15	< 0.001	-0.19	< 0.001
duration of GDP canceled	-0.055	0.033	-0.073	< 0.001

The reward regressors for the two airports are remarkably similar. This similarity is consistent with the conjecture that while different phenomena cause congestion issues leading to GDPs at these two airports (evidenced by the different important features in the two BC GDP initialization models described in sub-section IV.B.1), GDPs are implemented to achieve roughly the same objectives at both airports. Furthermore, this similarity illustrates the potential of IRL algorithms to identify reward functions and corresponding policies that generalize—that work in a variety of contexts, including those not represented in training data. More specifically, if non-constant reward features that achieve  $p$ -values less than 0.05 are sorted from largest to smallest magnitude of the average corresponding parameter estimate, the order of the top four features is identical for the two airports: the indicator that the air buffer at the end of the time step is greater than or equal to five, the number of unused slots during rate control, the duration of GDP canceled, and the change in the air buffer. These parameter estimates are all negative, as would be expected. The estimates also quantify the balance achieved by traffic managers as they face a fundamental trade-off in GDP implementation: airborne delay is expensive, and implementing GDPs can help achieve it, but excessive GDP implementation can lead to undesired under-utilization of available capacity. The objective functions used in various algorithms,<sup>2,24</sup> including one in an operational GDP decision-support tool,<sup>12</sup> all specify some balance for this trade-off. However, as far as we know, this is the first time that the balance achieved in current operations has been inferred directly from historical traffic flow management initiatives and related data.

If we investigate the average parameter estimates for less important features (those that achieve relatively high  $p$ -values and/or average parameter estimates with small magnitudes), then we find some differences between the regressors for the two airports. For example, while the average parameter estimates for buffer levels and changes in buffer levels are all negative for the regressors for SFO, as would be expected, this is not the case for EWR. Furthermore, while the average estimate of the parameter corresponding to arrivals is slightly positive for the SFO regressors, this value is slightly negative for the EWR regressors. These counter-intuitive parameter estimates may help explain why the EWR IRL GDP implemented model over-predicted GDP initialization.

As was mentioned in Section III.A.2, we selected discount factors of 0.4 for EWR and 0.2 for SFO based on how well the regressors were able to fit the resulting reward sample data (as measured by  $R^2$ ). This suggests that when determining whether or not to use a GDP, decision makers are concerned with rewards extending only a few hours into the future, and that by far the greatest importance is given to rewards in the current time step. This is consistent our observation that predictions of conditions more than four hours into the future do not improve the predictive performance of the BC GDP implemented models.

## V. Conclusions

GDPs seem to be a tool for strategically managing traffic in an effort to achieve desired values for certain metrics, suggesting that IRL may be a promising technique for GDP analytics. Therefore, we compared IRL models of GDP implementation to BC models. More precisely, we developed BC models of GDP implementation that are based on random forest models of GDP initialization and GDP cancellation. We used the CSI IRL algorithm to infer reward functions consistent with historical state and action data and then used rollouts to find policies attempting to optimize expected total discounted reward objectives based on the inferred reward functions. Furthermore, we implemented BC models for GDP parameters that are used by the rollouts policies. The models were developed for EWR and SFO and evaluated using cross validation on a data set consisting of 455 days in the summers of 2011–2013.

The BC GDP implemented models we developed for EWR and SFO demonstrate substantially stronger predictive performance than the IRL GDP implemented models we developed when predicting GDP implementation on testing data. However, our experiments also suggest that none of the models predict the rare GDP initialization or cancellation events well. Although they do not prove anything about the decision-making process, the results presented here suggest that decisions regarding whether or not to implement a GDP in a given hour are better modeled as a reaction to the current situation rather than an effort to deterministically and strategically achieve rewards accrued over time.

We also investigated the structure of the models in order to gain insights into GDP implementation behavior. Feature importance scores derived from the structure of the random forest BC GDP initialization and GDP cancellation models suggest that the set of most important features varies between airports; features related to scheduled arrivals, predicted airport arrival capacity levels, the previous GDP plan, certain weather conditions, and reroutes are most important for one or both airports. The reward functions inferred by the IRL algorithm are not able to achieve a good fit of the training data, but their structures suggest that decision makers at both airports are primarily concerned with avoiding relatively large numbers of flights that must incur delay in the air, avoiding unused arrival slots while delaying flights on the ground to achieve a certain rate of arrivals at the airport, and avoiding canceling a GDP long before its planned end time.

Features related to predictions of conditions more than four hours in the future do not improve the predictive power of the BC GDP implemented models. Similarly, we selected low discount factors of 0.4 for EWR and 0.2 for SFO for use in the IRL algorithm because these values led to reward training data that the reward regressor was better able to fit. These characteristics of the BC and IRL models are inferred from historical data and suggest that GDP implementation decisions are made primarily based on conditions now or conditions anticipated in the next couple of hours.

## VI. Future Work

The techniques used to train the many models involved in this research could be refined, and different types of models could be evaluated. For example, our intuition is that the ground and air buffer system model might be too simple, so it may be wise to utilize a more complex model that leverages higher-resolution traffic demand data (e.g. flight plans). Furthermore, we could investigate stochastic models of GDP decision making. These models might be more realistic in the sense that the same GDP decisions might not be made each time the system is in a particular state, given the data we use to define the state. Furthermore, probabilistic predictions might be more useful and insightful than deterministic predictions.

## Acknowledgments

We are grateful to Shon Grabbe, Banavar Sridhar, Avijit Mukherjee, Deepak Kulkarni, and Heather Arneson for providing feedback on this research as it progressed. We are also thankful for feedback on the initial plan for this research from Tony Evans, Roberto Bunge, Ryder Winck, Paul Varkey, Nikunj Oza, Karl Bilimoria, and Tatsuya Kotegawa. David Hattaway put us in touch with Cindy Hood, Supervisory Traffic Management Coordinator at New York TRACON, who we thank for valuable insights into current GDP decision making.

## References

- <sup>1</sup>Rios, J., "Aggregate Statistics of National Traffic Management Initiatives," *AIAA Aviation Technology, Integration, and Operations Conference*, Fort Worth, TX, October 2010.
- <sup>2</sup>Sridhar, B., Grabbe, S. R., and Mukherjee, A., "Modeling and Optimization in Traffic Flow Management," *Proceedings of the IEEE*, Vol. 96, No. 12, December 2008.
- <sup>3</sup>Grabbe, S., Sridhar, B., and Mukherjee, A., "Similar Days in the NAS: an Airport Perspective," *AIAA Aviation Technology, Integration, and Operations Conference*, September 2013.
- <sup>4</sup>Ratliff, N., Ziebart, B., Peterson, K., Bagnell, J. A., and Hebert, M., "Inverse Optimal Heuristic Control for Imitation Learning," Paper 48, Carnegie Mellon University Robotics Institute, 2009.
- <sup>5</sup>Wolfe, S. R. and Rios, J. L., "A Method for Using Historical Ground Delay Programs to Inform Day-of-Operations Programs," *AIAA Guidance, Navigation, and Control Conference*, Portland, OR, August 2011.
- <sup>6</sup>Wang, Y. and Kulkarni, D., "Modeling Weather Impact on Ground Delay Programs," *SAE Journal of Aerospace*, Vol. 4, No. 2, November 2011, pp. 1207–1215.
- <sup>7</sup>Bloem, M., Hattaway, D., and Bambos, N., "Evaluation of Algorithms for a Miles-in-Trail Decision Support Tool," *International Conference on Research in Air Transportation*, Berkeley, CA, May 2012.
- <sup>8</sup>Kulkarni, D., Wang, Y., and Sridhar, B., "Data Mining for Understanding and Improving Decision-Making Affecting Ground Delay Programs," *Proc. of AIAA/IEEE Digital Avionics Systems Conference*, Syracuse, NY, October 2013.
- <sup>9</sup>Mukherjee, A., Grabbe, S., and Sridhar, B., "Predicting Ground Delay Program At An Airport Based on Meteorological Conditions," *AIAA Aviation Technology, Integration, and Operations Conference*, Atlanta, GA, June 2014.
- <sup>10</sup>Ratliff, N. D., Bagnell, J. A., and Zinkevich, M. A., "Maximum Margin Planning," *Proc. of International Conference on Machine Learning*, Pittsburgh, PA, 2006.
- <sup>11</sup>Abbeel, P. and Ng, A. Y., "Apprenticeship Learning via Inverse Reinforcement Learning," *Proc. of International Conference on Machine Learning*, Banff, Canada, 2004.
- <sup>12</sup>Cook, L. S. and Wood, B., "A Model for Determining Ground Delay Program Parameters Using a Probabilistic Forecast of Stratus Clearing," *Proc. of 8th USA/Europe Air Traffic Management Research & Development Seminar*, Napa, CA, June 2009.
- <sup>13</sup>Liu, Y. and Hansen, M., "Ground Delay Program Decision-making using Multiple Criteria: A Single Airport Case," *USA/Europe Air Traffic Management Research & Development Seminar*, Chicago, IL, June 2013.
- <sup>14</sup>Federal Aviation Administration, "FAA Operations & Performance Data," <http://aspm.faa.gov/>.
- <sup>15</sup>Liu, P.-C., *Managing Uncertainty in the Single Airport Ground Holding Problem Using Scenario-based and Scenario-free Approaches*, PhD dissertation, University of California, Berkeley, CA, 2007.
- <sup>16</sup>Smith, D. A. and Sherry, L., "Decision Support Tool for Predicting Aircraft Arrival Rates, Ground Delay Programs, and Airport Delays from Weather Forecasts," *International Conference on Research in Air Transportation*, Fairfax, VA, February 2008.
- <sup>17</sup>Wang, Y., "Prediction of Weather Impacted Airport Capacity using Ensemble Learning," *AIAA/IEEE Digital Avionics Systems Conference*, Seattle, WA, October 2011.
- <sup>18</sup>Buxi, G. and Hansen, M., "Generating Probabilistic Capacity Profiles from weather forecast: A design-of-experiment approach," *USA/Europe Air Traffic Management Research & Development Seminar*, Berlin, Germany, June 2011.
- <sup>19</sup>Wang, Y., "Prediction of Weather Impacted Airport Capacity using RUC-2 Forecast," *AIAA/IEEE Digital Avionics Systems Conference*, Williamsburg, VA, October 2012.
- <sup>20</sup>Provan, C. A., Cook, L., and Cunningham, J., "A Probabilistic Airport Capacity Model for Improved Ground Delay Program Planning," *AIAA/IEEE Digital Avionics Systems Conference*, Seattle, WA, October 2011.
- <sup>21</sup>Cunningham, J., Cook, L., and Provan, C., "The Utilization of Current Forecast Products in a Probabilistic Airport Capacity Model," *AMS Annual Meeting*, New Orleans, LA, January 2012.
- <sup>22</sup>Dhal, R., Roy, S., Taylor, C., and Wanke, C., "Forecasting Weather-Impacted Airport Capacities for Flow Contingency Management: Advanced Methods and Integration," *AIAA Aviation Technology, Integration, and Operations Conference*, Los Angeles, CA, August 2013.
- <sup>23</sup>Kim, A. and Hansen, M., "Deconstructing delay: A non-parametric approach to analyzing delay changes in single server queuing systems," *Transportation Research Part B*, Vol. 58, December 2013, pp. 119–133.
- <sup>24</sup>Ball, M., Barnhart, C., Nemhauser, G., and Odoni, A., "Air Transportation: Irregular Operations and Control," *Handbooks in Operations Research & Management Science*, edited by C. Barnhart and G. Laporte, Vol. 14, chap. 1, Elsevier, 2007, pp. 1–61.
- <sup>25</sup>Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Édouard Duchesnay, "Scikit-learn: Machine Learning in Python," *Journal of Machine Learning Research*, Vol. 12, 2011, pp. 2825–2830.
- <sup>26</sup>Hastie, T., Tibshirani, R., and Friedman, J., *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer-Verlag, New York, 2001.
- <sup>27</sup>Chawla, N. V., Bowyer, K. W., Hall, L. O., and Kegelmeyer, W. P., "SMOTE: Synthetic Minority Over-sampling Technique," *Journal of Artificial Intelligence Research*, Vol. 16, 2002, pp. 321–357.
- <sup>28</sup>Jeschkies, K., "SMOTE implementation for over-sampling," <http://comments.gmane.org/gmane.comp.python.scikit-learn/5278>, November 2012.
- <sup>29</sup>Klein, E., Piot, B., Geist, M., and Pietquin, O., "A Cascaded Supervised Learning Approach to Inverse Reinforcement Learning," *Machine Learning and Knowledge Discovery in Databases*, edited by H. Blockeel, K. Kersting, S. Nijssen, and F. Železný, Vol. 8188 of *Lecture Notes in Computer Science*, Springer, September 2013, pp. 1–16.
- <sup>30</sup>Bertsekas, D. P., *Dynamic Programming and Optimal Control*, Vol. 1, Athena Scientific, Nashua, NH, 2005.
- <sup>31</sup>Ng, A., "Advice for applying Machine Learning," <http://cs229.stanford.edu/materials/ML-advice.pdf>, 2011.